Online submodular welfare maximization: Greedy is optimal

Michael Kapralov Ian Post Jan Vondrák

MIT, Waterloo, IBM Almaden

July 31, 2013
FND Workshop

Welfare Maximization

- A set M of $m = |M|$ items
- n agents each with valuations $w_i(S)$ for sets of items $S \subseteq M$
- Goal: assign items to agents to maximize $\sum_i w_i(S_i)$

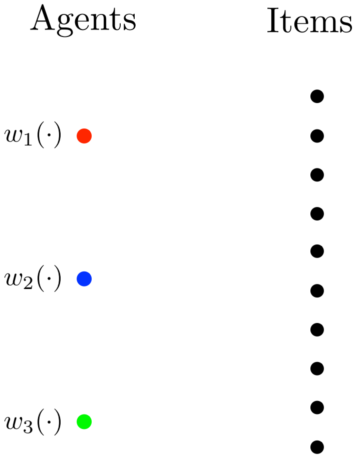

Welfare Maximization

- A set M of $m = |M|$ items
- n agents each with valuations $w_i(S)$ for sets of items $S \subseteq M$
- Goal: assign items to agents to maximize $\sum_i w_i(S_i)$

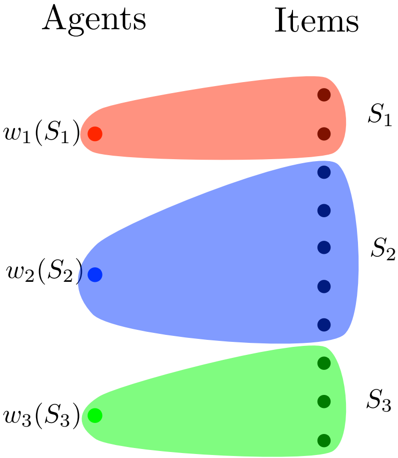

Welfare Maximization

- A set M of $m = |M|$ items
- n agents each with valuations $w_i(S)$ for sets of items $S \subseteq M$
- Goal: assign items to agents to maximize $\sum_i w_i(S_i)$
- Items arrive online and must be irrevocably assigned

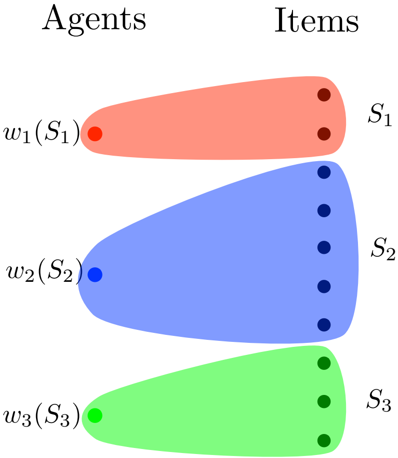

Submodularity

We will assume valuations $w_i(\cdot)$ are *monotone* and *submodular*

Definition

A valuation function $w : 2^M \rightarrow \mathbb{R}_+$ *is* submodular *if* $w(T + j) - w(T) \leq w(S + j) - w(S)$ *for all* $j \in M$ *whenever* $S \subseteq T$*, and* monotone *if* $w(S) \leq w(T)$ *whenever* $S \subseteq T$.

Submodularity

Definition

A valuation function $w : 2^M \to \mathbb{R}_+$ *is a* coverage valuation *if there is a set system* $\{A_j : j \in M\}$ *such that* $w(S) = |\bigcup_{j \in S} A_j|$ *for all* $S \subseteq M$*.*

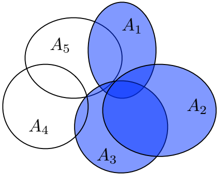

Definition

A function $w : 2^M \to \mathbb{R}_+$ *is* budget-additive *if* $w(S) = \min\{B, \sum_{j \in S} b_j\}$ *for some budget* B *and item values* b_j*.*

Previous Work

- The greedy algorithm is a $1/2$-approximation for monotone submodular valuations [FNW78, LLN06]
- A $(1 - 1/e)$-approximation is known offline [Von08], and this is optimal [KLMM08]
- $(1 - 1/e)$-approximations are known online for several special cases
  - Matching [KVV90]
  - Budget-additive with small bids ($b_{ij} \ll B_i$) [MSVV05]
  - Budget-additive with single bids ($b_{ij} \in \{0, b_i\}$ for some b_i) [AGKM11]
  - Conjectured for general budget-additive valuations
- In some stochastic settings it is possible to beat $1 - 1/e$ [GM08, FMMM09, BK10, MOS11, MOZ12]

Is $1 - 1/e$ possible online with general monotone submodular valuations?

Main Result

No, $1/2$ is optimal:

Theorem

For any constant $\delta > 0$, there is no $(1/2 + \delta)$-competitive polynomial-time algorithm (even randomized, against an oblivious adversary) for the online welfare maximization problem with coverage valuations unless $NP = RP$.

Main Result

No, $1/2$ is optimal:

Theorem

For any constant $\delta > 0$, there is no $(1/2 + \delta)$-competitive polynomial-time algorithm (even randomized, against an oblivious adversary) for the online welfare maximization problem with coverage valuations unless $NP = RP$.

- Relies on a careful combination of two sources of hardness, both computational (inapproximability of Max k-cover [Fei98]) and information-theoretic (the unknown online ordering)

Other Results

Theorem

For budget-additive valuations, no online (randomized) algorithm can achieve (in expectation) more than a 0.612-fraction of the optimal value of the standard LP relaxation.

Other Results

Theorem

For budget-additive valuations, no online (randomized) algorithm can achieve (in expectation) more than a 0.612-fraction of the optimal value of the standard LP relaxation.

- Does not rule out a $(1 - 1/e)$-approximation for budget-additive valuations

Other Results

Theorem

In the stochastic i.i.d. model the greedy algorithm is $(1 - 1/e)$*-competitive for valuations satisfying the property of* diminishing returns, *and no polynomial-time algorithm can achieve* $(1 - 1/e + \delta)$ *unless* $NP = RP$.

Other Results

Theorem

In the stochastic i.i.d. model the greedy algorithm is $(1 - 1/e)$-competitive for valuations satisfying the property of diminishing returns, *and no polynomial-time algorithm can achieve $(1 - 1/e + \delta)$ unless $NP = RP$.*

- *Diminishing returns* is a natural generalization of submodularity to multisets that we define. We believe that it may be useful in other problems.

Hardness of Max k-cover

Theorem (Fei98, FT04, FV10)

For any fixed $c > 0$ and $\epsilon > 0$, it is NP-hard to distinguish between the following two cases for a given collection of sets $\mathcal{S} \subset 2^U$, partitioned into groups $\mathcal{S}_1, \ldots, \mathcal{S}_k$:

- *YES case: There are k disjoint sets, 1 from each group $\mathcal{S}_i$, whose union is the universe U.*
- *NO case: For any choice of $\ell \leq ck$ sets, their union covers at most a $(1 - (1 - 1/k)^\ell + \epsilon)$-fraction ($\approx (1 - e^{-\ell/k} + \epsilon)$) of the elements of U.*

This holds even for set systems such that

- *every set has the same (constant) size s; and*
- *each group contains the same (constant) number of sets.*

Hardness of offline welfare maximization

- Previously proved by [KLMM08] via a different technique
- Consider a hard instance of Max k-cover where n is the number of sets in each group and k is the number of groups
- Create n agents and $m = kn$ items
- For agent i the item $(j_1, j_2) \in [k] \times [n]$ is associated with the set $A_{j_1, j_2 + i \pmod n}$, so the value of a set S of items is

$$w_i(S) = \left| \bigcup_{(j_1, j_2) \in S} A_{j_1, (j_2 + i \bmod n)} \right| .$$

- YES instance: set $S_i = \{(j, \pi(j) - i \bmod n)\}$ where $\pi : [k] \to [n]$ is the set from each group covering U, which gives $w_i(S_i) = |U|$ for all i
- NO instance: $(1 - e^{-\ell/k} + \epsilon)|U|$ is concave, so maximized when each agent gets k items, which yields $w_i(S_i) \le (1 - 1/e + \epsilon)|U|$.

Hardness of online matching

Lower bound for online matching [KVV90]:

- At each step a random agent “drops out” and cannot be matched to any future element
- If there are t steps remaining, we can only guess the identity of this agent with probability $1/t$
- Leads to a $1 - 1/e$ lower bound

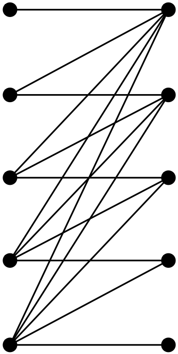

Hardness of online matching

Lower bound for online matching [KVV90]:

- At each step a random agent "drops out" and cannot be matched to any future element
- If there are t steps remaining, we can only guess the identity of this agent with probability $1/t$
- Leads to a $1 - 1/e$ lower bound

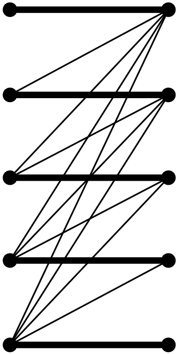

Hardness of online matching

Lower bound for online matching [KVV90]:

- At each step a random agent "drops out" and cannot be matched to any future element
- If there are t steps remaining, we can only guess the identity of this agent with probability $1/t$
- Leads to a $1 - 1/e$ lower bound

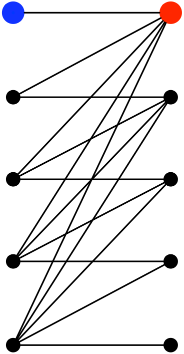

Hardness of online matching

Lower bound for online matching [KVV90]:

- At each step a random agent "drops out" and cannot be matched to any future element
- If there are t steps remaining, we can only guess the identity of this agent with probability $1/t$
- Leads to a $1 - 1/e$ lower bound

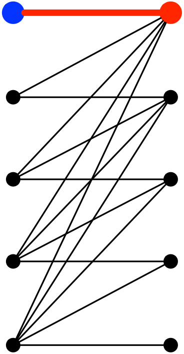

Hard instance of online welfare maximization

Idea: expand each vertex into an entire instance of welfare maximization with coverage valuations to impose the additional difficulty of approximating an APX-hard problem at each stage

- tn agents and tm items
- At each stage 1 copy of each item arrives, and a random copy of each agent drops out

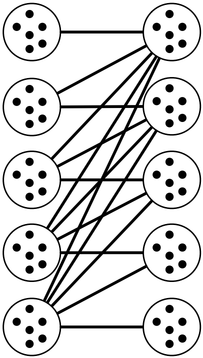

Hard instance of online welfare maximization

Idea: expand each vertex into an entire instance of welfare maximization with coverage valuations to impose the additional difficulty of approximating an APX-hard problem at each stage

- *tn* agents and *tm* items
- At each stage 1 copy of each item arrives, and a random copy of each agent drops out

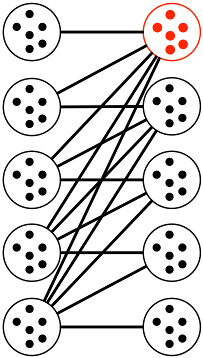

Hard instance of online welfare maximization

Idea: expand each vertex into an entire instance of welfare maximization with coverage valuations to impose the additional difficulty of approximating an APX-hard problem at each stage

- *tn* agents and *tm* items
- At each stage 1 copy of each item arrives, and a random copy of each agent drops out

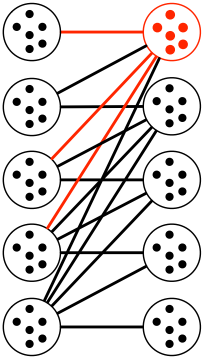

Hard instance of online welfare maximization

Idea: expand each vertex into an entire instance of welfare maximization with coverage valuations to impose the additional difficulty of approximating an APX-hard problem at each stage

- tn agents and tm items
- At each stage 1 copy of each item arrives, and a random copy of each agent drops out

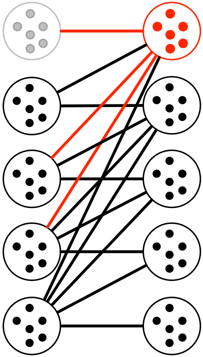

Difficulty

Problem: items in one stage may be spread around agents in many different copies

- But we have a bound depending only on the number of items allocated to any agent:
  - In the NO instance, any choice of $\ell \leq ck$ sets covers at most $(1 - e^{-\ell/k} + \epsilon)$-fraction of U
  - Taking multiple copies of the same item does not increase valuation

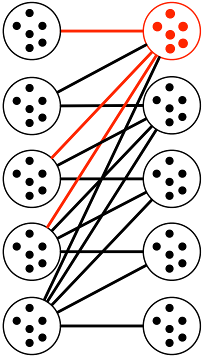

Proof sketch

Lemma

For all $\epsilon' > 0$ *there exist* $\epsilon, c > 0$ *and a constant lower bound on* k *such that the expected value collected in the NO case by the agents deactivated at the end of stage* j *is at most* $(j/t + \epsilon')n|U|$ *for* $j \leq (1 - \epsilon')t$.

Sketch.

- Because the agents deactivated in each stage are random, we can bound the expected number of items allocated to agents deactivated at the end of stage j by $m \ln \frac{t}{t-j}$
- An agent receiving ℓ items gets about $(1 - e^{-\ell/k})|U|$ value
- This is concave and maximized if each agent gets $\ell = k \ln \frac{t}{t-j}$ items:

$$\left(1 - e^{(k \ln \frac{t}{t-j})/k}\right)|U| = \frac{j}{t}|U|$$

□

Proof sketch

Sketch of Theorem.

- By the previous lemma in the NO instance we achieve at most

$$\sum_{j=1}^{(1-\epsilon')t} \left(\frac{j}{t} + \epsilon'\right) n|U| + \epsilon' tn|U| \leq \left(\frac{1}{2} + 2\epsilon'\right) tn|U| \ .$$

- Suppose we have a $(1/2 + \delta)$-approximation. Set $\epsilon' = \delta/4$, then in the YES instance it achieves at least

$$\left(\frac{1}{2} + 4\epsilon'\right) tn|U| \ .$$

- We can distinguish between these two cases with constant probability, and this can be improved to one-sided error

□

Budget-additive valuations

$$\begin{array}{llrl} \text{maximize} & & \sum_{i\in N, j\in M} b_{ij}x_{ij} & \\ \text{subject to} & \forall i \in N, & \sum_{j\in M} b_{ij}x_{ij} & \leq B_i \\ & \forall j \in M, & \sum_{i\in N} x_{ij} & \leq 1 \\ & & x & \geq 0 \end{array}$$

Theorem

No online (randomized) algorithm can achieve (in expectation) more than a 0.612-fraction ($< 1 - 1/e$) of the optimal value of this LP.

- Similar proof: take a particular instance with an integrality gap and make it hard to approximate online by blowing it up into an instance with the same structure as the hard instance of online matching.

Submodularity on multisets

Recall that $f : 2^M \to \mathbb{R}_+$ is *submodular* if
$f(S+j) - f(S) \geq w(T+j) - w(T)$ for all $j \in M$ whenever $S \subseteq T$

Definition

A function $f : \mathbb{Z}_+^m \to \mathbb{R}$ *has the property of* diminishing returns, *if for any* $x \leq y$ *(coordinate-wise) and any unit basis vector* $e_i = (0, \ldots, 0, 1, 0, \ldots, 0)$*,* $i \in [m]$*,*

$$f(x + e_i) - f(x) \geq f(y + e_i) - f(y) \; .$$

- Note this implies f is submodular if its domain is $\{0, 1\}^m$ and concave if its domain is $\mathbb{Z}_+$

Stochastic welfare maximization

Theorem

Suppose the items are i.i.d. samples from some (possibly unknown) distribution over a set M. Then the greedy algorithm achieves a $(1 - 1/e)$-approximation, and there is no $(1 - 1/e + \delta)$-approximation unless $NP = RP$.

- The upper bound is a straightforward extension of the analysis of [DJSW11] for the budget-additive case
  - If the current allocation is $(T_1, \ldots, T_n)$, then the expected gain of the next item is at least $\frac{1}{m}(LP - \sum_i w_i(T_i))$
- For the lower bound, take the above $(1 - 1/e)$-hard instance of offline welfare maximization, make $t = \text{poly}(m)$ copies of each agent, and draw tm random samples from the set of items

Open Questions

- Is $1 - 1/e$ possible online for budget-additive functions?
  - If you think yes:
    - Configuration LP
    - Online matching without *monotonicity*
  - If you think no:
    - Probably requires stronger hardness result for offline version
- Is there an offline $(1 - 1/e)$-approximation for functions with diminishing returns?